arXiv:1608.02413v2 [cs.DS] 17 Nov 2016

EPR-dictionaries: A practical and fast data structure for constant time searches in unidirectional and bidirectional FM-indices

Christopher Pockrandt*, Marcel Ehrhardt**, and Knut Reinert***

Department of Computer Science and Mathematics, Freie Universität Berlin, Germany

Abstract. The unidirectional FM index was introduced by Ferragina and Manzini in 2000 and allows to search a pattern in the index in one direction. The bidirectional FM index (2FM) was introduced by Lam et al. in 2009. It allows to search for a pattern by extending an infix of the pattern arbitrarily to the left or right. The method of Lam et al. can conduct one step in time $\mathcal{O}(\sigma)$ while needing space $\mathcal{O}(\sigma \cdot n)$ using constant time rank queries on bit vectors. Schnattinger and colleagues improved this time to $\mathcal{O}(\log \sigma)$ while using $\mathcal{O}(\log \sigma \cdot n)$ bits of space for both, the FM and 2FM index. This is achieved by the use of binary wavelet trees.

In this paper we introduce a new, practical method for conducting an exact search in a uni- and bidirectional FM index in $\mathcal{O}(1)$ time per step while using $\mathcal{O}(\log \sigma \cdot n) + o(\log \sigma \cdot \sigma \cdot n)$ bits of space. This is done by replacing the binary wavelet tree by a new data structure, the *Enhanced Prefixsum Rank dictionary* (EPR-dictionary).

We implemented this method in the SeqAn C++ library and experimentally validated our theoretical results. In addition we compared our implementation with other freely available implementations of bidirectional indices and show that we are between $\approx 2.6 - 4.8$ times faster. This will have a large impact for many bioinformatics applications that rely on practical implementations of (2)FM indices e.g. for read mapping. To our knowledge this is the first implementation of a constant time method for a search step in 2FM indices.

Keywords: FM index, bidirectional, BWT, bit vector, rank queries, read mapping.

1 Introduction

It is seldom that new data structures or algorithms have such a large practical impact as full text indices had for biological sequence analysis. The so-called next-generation sequencing (NGS) allows to produce billions of small DNA strings called *reads*, usually of length 100-250. It is an invaluable technology for a multitude of applications in biomedicine. In many of these applications finding the positions of the DNA strings in a reference genome (i.e., a large string of length $\approx 10^7 - 10^{10}$) is the first fundamental step preceding downstream analyses. Finding the positions of the reads is commonly referred to as *read mapping*.

* FU Berlin (christopher.pockrandt@fu-berlin.de, http://reinert-lab.de).
** FU Berlin (marcel.ehrhardt@fu-berlin.de).
*** FU Berlin (knut.reinert@fu-berlin.de)

Because of sequencing errors and genomic variations not all strings occur exactly in a reference genome. Therefore approximate occurrences must be considered and algorithms for approximate string matching tolerating mismatches, insertions, and deletions must be applied to solve the problem.

This has triggered a plethora of work in the field to implement fast and accurate read mappers. Many of the popular programs like Bowtie2 [12], BWA [14], BWA-Mem [13], Masai [20], Yara [19], and GEM [17] use as their main data structure a version of the FM index [5] that was introduced by Ferragina and Manzini in 2000. The FM index is based on the Burrows-Wheeler transform (BWT) [3] of the given text, i.e., the genomes at hand, and conceptually some lookup tables containing counts of characters in prefixes of the text. In its original form it allows to search exactly for a pattern in *one direction* by matching the characters of the pattern with characters in the BWT [3] (i.e., extending a suffix of the pattern character by character to the left). It was later extended to the 2FM index by Lam et al. [10] and Schnattinger et al. [18]. The 2FM index allows to search in both directions, that means we can extend an infix of a pattern arbitrarily to the left or to the right. In order to reduce its space requirements, the count tables are in practice replaced by efficient bit vector data structures with rank support (see for example [9]). The search method of Lam et al. can conduct one search step in a 2FM index in time $\mathcal{O}(\sigma)$ while needing space $\mathcal{O}(\sigma \cdot n)$ using constant time rank queries on bit vectors. Schnattinger et al. improved this time to $\mathcal{O}(\log \sigma)$ while using $\mathcal{O}(\log \sigma \cdot n)$ bits of space for both, the FM and 2FM index. This is achieved by the use of binary wavelet trees introduced by Grossi et al. [8]. In the last years several theoretical results appeared that improved on this. However, none of those has found a way into a practical implementation.

In this paper we introduce a new method for conducting an exact search in a uni- and bidirectional FM index that needs $\mathcal{O}(1)$ time per step while using $\mathcal{O}(\log \sigma \cdot n) + o(\log \sigma \cdot \sigma \cdot n)$ bits of space. This is done by replacing the binary wavelet tree by a new data structure, the *Enhanced Prefixsum Rank dictionary* (EPR-dictionary). To our knowledge this is the first implementation of a constant time method for 2FM indices. We will show, that the method outperforms other implementations by several factors at the expense of a slight increase in memory usage resulting in a very practical method.

In the following paragraph we will review the concepts of the FM and 2FM index as well as constant time rank queries very shortly (readers unfamiliar with this can find a more detailed description in the appendix).

1.1 Introduction to the FM and 2FM Index

Given a text T of length n over an ordered, finite alphabet $\Sigma = \{c_1, \ldots, c_\sigma\}$ with $\forall 1 \leq i < \sigma : c_i <_{lex} c_{i+1}$, let $T[i]$ denote the character at position i in T, $\cdot$ the concatenation operator and $T[1..i]$ the prefix of T up to the character at position i. T^{rev} represents the reversed text. We assume that T ends with a sentinel character $\$ \notin \Sigma$ that does not occur in any other position in T and is lexicographically smaller than any character in Σ. The FM index needs the Burrows-Wheeler transform (BWT) of T. The BWT is the concatenation of characters in the last column of all lexicographically sorted cyclic permutations of the string T. We will refer to the BWT as L.

In contrast to suffix trees or suffix arrays, where a prefix P of a pattern is extended by characters to the right (referred to as forward search $P \rightarrow Pc$ for $c \in \Sigma$), the FM

index can only be searched using backward searches, i.e., extending a suffix P' by characters to the left, $P' \rightarrow cP'$. Performing a single character backward search of c in the FM index will require two pieces of information. First, $C(c)$, the number of characters in L that are lexicographically smaller than c, second, $Occ(c, i)$, the number of c's in $L[1..i]$. Given a range $[a, b]$ for P; i.e., the range in the sorted list of cyclic permutations that start with P, we can compute the range $[a', b']$ for cP as follows: $[a', b'] = [C(c) + Occ(c, a - 1) + 1, C(c) + Occ(c, b)]$.

The 2FM index maintains two FM indices $\mathcal{I}$ and $\mathcal{I}^{rev}$, one for the original text T and one for the reversed text T^{rev}. Searching a pattern left to right on the original text (i.e., conducting a forward search) corresponds to a backward search in $\mathcal{I}^{rev}$; searching a pattern right to left in the original text corresponds to a backward search in $\mathcal{I}$. The difficulty is to keep both indices *synchronized* whenever a search step is performed. W.l.o.g. we assume that we want to extend the pattern to the right, i.e., perform a forward search $P \rightarrow Pc_j$ for some character c_j. First, the backward search $P^{rev} \rightarrow c_j P^{rev}$ is carried out on $\mathcal{I}^{rev}$ and its new range $[a'_{rev}, b'_{rev}] = [C(c_j) + Occ(c_j, a_{rev} - 1) + 1, C(c_j) + Occ(c_j, b_{rev})]$ is computed. The new range in $\mathcal{I}$ can be calculated using the interval $[a, b]$ for P in $\mathcal{I}$ and the range size of the reversed texts index $[a', b'] = [a + smaller, a + smaller + b'_{rev} - a'_{rev}]$. To compute $smaller$, e.g. Lam et al. [10] propose to perform $\mathcal{O}(\sigma)$ backward searches $P^{rev} \rightarrow c_i P^{rev}$ for all $1 \leq i < j$ and sum up the range sizes, i.e., $smaller = \sum_{1 \leq i < j} Occ(c_i, b_{rev}) - \sum_{1 \leq i < j} Occ(c_i, a_{rev} - 1)$ leading to a total running time of $\mathcal{O}(\sigma)$.

The implementation of the occurrence table Occ is usually *not* done by storing explicitly the values of the entire table. Instead of storing the entire $Occ : \Sigma \times \{1, \ldots, n\} \rightarrow \{1, \ldots, n\}$ one uses the more space-efficient *constant time rank dictionary*: for every $c \in \Sigma$ a bit vector $B_c[1..n]$ is constructed such that $B_c[i] = 1$ if and only if $L[i] = c$. Thus the occurrence value equals the number of 1's in $B_c[1..i]$, i.e., $Occ(c, i) = rank(B_c, i)$. Jacobson [9] showed that rank queries can be answered in constant time using only $o(n)$ additional space per bit vector by employing a sum of two count arrays (i.e., *blocks* and *superblocks*) and a final *in-block* count. Since then many other constant time rank query data structures have been proposed. For an overview we refer the reader to [16] containing a comparison of various implementations. For readers unfamiliar with *2-level rank dictionaries*, an explanation is given in the appendix.

1.2 Recent improvements on the FM and 2FM index

For large alphabets, it is not practical to maintain for each character a bit vector with rank support. In 2003 Grossi et al. [8] proposed the use of a more space efficient data structure for the FM index, called the *(binary) wavelet tree* (WT) that was later used by Schnattinger [18] for an implementation of bidirectional FM indices. It is a binary tree of height $\mathcal{O}(\log \sigma)$ with a bit vector of length n with rank support at each level. This reduces the space consumption by a factor of $\mathcal{O}(\frac{\log \sigma}{\sigma})$ in trade-off for an increased running time of $\mathcal{O}(\log \sigma)$. Schnattinger used the fact that not only the rank query for a given character c can be computed in $\mathcal{O}(\log \sigma)$ but also the $smaller$ value can be computed in the same asymptotic time which is quite convenient for the 2FM index. Ferragina et al. proposed a new data structure in 2007 [6], the multi-ary wavelet tree, which could be used to speed up the needed rank queries of 2FM indices. In 2013 Belazzougui et al. proposed the first constant-time bidirectional FM index [1]

using minimal perfect hashing, of which to our knowledge no implementation exists (see also [2] for an extended version). Our solution is based on bit vectors with rank support, which proved so far to be very fast in practice, in particular due to the `popcount` machine operation.

2 Theoretical results

In this section we present the main results of this paper. They are based on a simple observation and a new bit vector data structure with rank query support which allows us to improve upon the results of Lam and Schnattinger. Our proposed method runs in constant time per step while using $\mathcal{O}(\log \sigma \cdot n) + o(\log \sigma \cdot \sigma \cdot n)$ bits of space for small alphabets (i.e., $\sigma < \log(n)/\log\log(n)$) which is in theory inferior in space consumption to the results of mentioned above (see [1]), but in practice very fast, and presents to our knowledge the first constant time implementation of 2FM indices with this space complexity.

Our first observation is simple. Instead of defining a bit vector for each $c \in \Sigma$ to map characters equal to c in L to 1's, we suggest using prefix sum bit vectors PB_c, i.e., $PB_c[i] = 1$ if and only if $L[i] \leq_{lex} c$ for $c \in \Sigma$.

Theorem 1. *A step in a bidirectional search can be performed in time $\mathcal{O}(1)$ using $\mathcal{O}(\sigma \cdot n)$ bits of space.*

Proof. We define *Prefix-Occ*$(c_j, i) = rank(PB_{c_j}, i)$; that means it counts the number of occurrences of a character lexicographically smaller or equal than c_j up to position i. *Prefix-Occ*(c_j, i) and thus the *smaller* value for the 2FM index can now be computed by a single rank query $rank(PB_{c_j}, i)$, the original $Occ(c_j, i)$ value for backward searches needs only two rank queries and a subtraction, namely $Occ(c_j, i) = rank(PB_{c_j}, i) - rank(PB_{c_{j-1}}, i)$ (for the lexicographically smallest character c_0 no subtraction is necessary).

Note that the bit vector for the lexicographically largest character can be omitted, since all bits will be set to 1 and thus $rank(PB_{c_\sigma}, i) = i$, $\forall\, 1 \leq i \leq n$.

Our next idea is the main result of this work and will allow us to reduce the space complexity for both the FM and the 2FM index while maintaining the optimal running time of $\mathcal{O}(1)$ per search step. Instead of using normal bit vectors we use directly the binary encoding of the BWT (an idea already used by BWT-SW[11]). We call our data structure *EPR-dictionary*, short for *Enhanced Prefixsum Rank dictionary*.

2.1 The EPR-dictionary

The general idea of the EPR-dictionary is as follows. Assuming an ordered alphabet $\Sigma = \{c_1, \ldots, c_\sigma\}$, each character c_i is encoded by the binary value $ord_2(c_i)$ of its rank i. Conceptually, we use the binary representation of the BWT to derive from it a *spaced* bit vector representation for PB_c for each $s \in \Sigma$. Then we compute the auxiliary data structures (i.e., blocks and superblocks (see also appendix)) for each of those vectors. After we have those auxiliary structures which only need $o(n)$ bits space, we *delete* the bit vectors and only retain the BWT. In practice the blocks and superblocks are computed directly by a linear scan on the BWT. For the last in-block

query, we show how to derive the counts *Prefix-Occ*(c_j, i) from the BWT in constant time using a number of logical and arithmetic operations.

W.l.o.g. we assume that the block length is an even multiple of $\log \sigma$ to avoid case distinctions in the proof. In practice this holds since the in-block query is performed with popcounts on registers the length of which is a power of 2. All bitmasks used for computing the in-block rank are exactly as long as a block. For a character c_j we define the *rank bitmask* $rb(c_j)$ to be a binary sequence of concatenations of the pattern $0^{\lceil \log \sigma \rceil - 1} 1 \cdot ord_2(c_j)$, i.e., $\lceil \log \sigma \rceil - 1$ many bits set to 0 followed by a 1 followed by the binary encoding of the character c_j. For the DNA alphabet with $\Sigma = \{A, C, G, T\}$ and its binary encodings $\{00, 01, 10, 11\}$ the rank bitmask for $G \in \Sigma$ is for example $rb(G) = 01 \cdot 10 \cdot 01 \cdot 10 \ldots$.

Counting the characters inside a block is done in two steps. The characters at even and odd positions are counted separately to generate space for an overflow bit. Therefore we need a bitmask M_E masking characters at even positions from the bit vector. M_E has 1s for each even block of length $\lceil \log \sigma \rceil$, i.e., $M_E = 00 \cdot 11 \cdot 00 \cdot 11 \cdot 00 \cdot 11 \ldots$. Finally, we need a bitmask BM which filters out the lowest bit of each odd $\log \sigma$ block, i.e., $BM = 01 \cdot 00 \cdot 01 \cdot 00 \cdot 01 \cdot 00 \ldots$ for $\sigma = 4$.

Step 1. We first take the characters at odd positions inside the corresponding block of the BWT, subtract it from $rb(c_i)$, which will result in the rightmost bit of even character positions to be set to 1 if and only if the character to the right is smaller or equal to c_i. We then obtain exactly those bits by masking with BM.

$$B_{EPR}(c_i)_E = (rb(c_i) - (BWT \& M_E)) \& BM$$

Step 2. We then take the characters at even positions inside the corresponding block of the BWT by shifting them $\lceil \log \sigma \rceil$ bits to the right and masking with M_E. We can now continue as in step 1 by subtracting it from $rb(c_i)$, which will again result in a 1 bit in the rightmost bit of even character positions to be set to 1 if and only if the character to the right is smaller or equal to c_i. We then apply the bitmask BM to filter only these rightmost bits.

$$B_{EPR}(c_i)_O = (rb(c_i) - ((\gg_{\lceil \log \sigma \rceil} BWT) \& M_E)) \& BM$$

Finally both bit vectors are merged with one of them shifted by 1 to the left avoiding the rightmost bits to overlap. In practice this is faster than two popcount operations.

$$B_{EPR}(c_i) = B_{EPR}(c_i)_E | (\ll_1 B_{EPR}(c_i)_O)$$

Since we used the binary encoding of the BWT, note that the underlying rank queries have to be adapted to *Prefix-Occ*$(c_j, i) = rank(B_{EPR}(c_j), (i-1) \cdot \lceil \log \sigma \rceil + 1)$. It follows directly that $Occ(c_j, i)$ can by computed in constant time by observing that

$$Occ(c_j, i) = \begin{cases} \textit{Prefix-Occ}(c_j, i) - \textit{Prefix-Occ}(c_{j-1}, i) & \text{if } j > 0 \\ \textit{Prefix-Occ}(c_j, i) & \text{otherwise} \end{cases}$$

The EPR-transformed bit vector $B_{EPR}(c_j)$ is now a "normal" bit vector and thus allows us to compute the prefix sums for a string in constant time. This improves the running time of the 2FM index and makes it optimal in terms of speed.

$rb(G)$		01	10	01	10	01	10	01	10
$BWT\&M_E$	–	00	01	00	01	00	11	00	11
		-	(C)	-	(C)	-	(T)	-	(T)
	=	01	01	01	01	00	11	00	11
BM	&	01	00	01	00	01	00	01	00
$B_{EPR}(G)_E$	=	01	00	01	00	00	00	00	00

(a) step 1

$rb(G)$		01	10	01	10	01	10	01	10
$(\gg_{\lceil \log \sigma \rceil} BWT)\&M_E$	–	00	00	00	10	00	10	00	00
		-	(A)	-	(G)	-	(G)	-	(A)
	=	01	10	01	00	01	00	01	10
BM	&	01	00	01	00	01	00	01	00
$B_{EPR}(G)_O$	=	01	00	01	00	01	00	01	00

(b) step 2

$B_{EPR}(G)_E$		01	00	01	00	00	00	00	00
$\ll_1 B_{EPR}(G)_O$	|	10	00	10	00	10	00	10	00
$B_{EPR}(G)$	=	11	00	11	00	10	00	10	00
popcount	=	6							

(c) retrieving $B_{EPR}(G)$

Fig. 1: An example for $\Sigma = \{A, C, G, T\}$ that shows how to perform an in-block rank query for characters smaller or equal to G of the BWT substring $ACGCGTAT$. The resulting bit vector $B_{EPR}(G)$ has a 1 for each character smaller or equal to G., i.e., all positions except those with a T.

Let us now take a look at the space consumption. For our exposition we define the block length of $\ell = \left\lfloor \frac{\log n}{2} \right\rfloor$ (if ℓ is not a multiple of $\lceil \log \sigma \rceil$ padding strategies can be applied). Given a $B_{EPR}(c_j)$, for the m-th superblock we count the number of 1's (i.e., the number of occurrences of characters smaller or equal to c_j in the corresponding BWT) from the beginning of B_{EPR} to the end of the superblock in $M'[m][j] = rank(B_{EPR}(c_j), m \cdot \ell^2)$. As there are $\left\lfloor \frac{\lceil \log \sigma \rceil \cdot n}{\ell^2} \right\rfloor$ superblocks and σ characters, M' can be stored in

$$\mathcal{O}\left(\sigma \cdot \frac{\log \sigma \cdot n}{\ell^2} \cdot \log n\right) = \mathcal{O}\left(\sigma \cdot \log \sigma \cdot \frac{n}{\log n}\right) = o(\sigma \log \sigma \cdot n)$$

bits. For the m-th block we count the number of 1's from the beginning of the overlapping superblock to the end of the block in $M[m][j] = rank\big(B_{EPR}[1 + k\ell..n](c_j), (m - k)\ell\big)$ where $k = \left\lfloor \frac{m-1}{\ell} \right\rfloor \ell$ is the number of blocks left of the overlapping superblock. M has $\left\lfloor \frac{\lceil \log \sigma \rceil \cdot n}{\ell} \right\rfloor$ entries for every character and can be stored in

$$\mathcal{O}\left(\sigma \cdot \frac{\log \sigma \cdot n}{\ell} \cdot \log \ell^2\right) = \mathcal{O}\left(\sigma \cdot \log \sigma \cdot \frac{n \cdot \log \log n}{\log n}\right) = o(\sigma \log \sigma \cdot n)$$

bits.

Let P be a precomputed lookup table such that for each possible infix V of a bit vector $B_{EPR}(c_j)$ of length ℓ, $i \in \left[1.. \left\lfloor \frac{\ell}{\log \sigma} \right\rfloor\right]$ and $c_j \in \Sigma$ holds $P[V][i] = rank(V, (i -$

$1) \cdot \lceil \log \sigma \rceil + 1)$. There are $2^\ell \cdot \left\lfloor \frac{\ell}{\log \sigma} \right\rfloor$ entries of value at most $\left\lfloor \frac{\ell}{\log \sigma} \right\rfloor$ and thus can be stored in

$$\mathcal{O}\left(2^\ell \cdot \frac{\ell}{\log \sigma} \cdot \log \frac{\ell}{\log \sigma}\right) = \mathcal{O}\left(2^{\frac{\log n}{2}} \cdot \log(n-\sigma) \cdot \log\log(n-\sigma)\right) =$$

$$\mathcal{O}\left(\sqrt{n} \cdot \log n \cdot \log\log n\right) = o(n)$$

bits. Note that we do need this lookup table only *once*, since the position and counting of the bits set to 1 is the same for all characters.

Equivalent to Theorem 1, we do not need to store blocks and superblocks for c_σ since $rank(B_{EPR}(c_\sigma), i) = i,\ \forall\, 1 \leq i \leq n$.

Theorem 2 (Constant time prefix sum query). *One search step in an FM index or 2FM index can be performed in $\mathcal{O}(1)$ time using $O(\log \sigma \cdot n) + o(\log \sigma \cdot \sigma \cdot n)$ bits of space.*

Proof. The BWT can be stored in $\mathcal{O}(\log \sigma \cdot n)$, all the blocks, superblocks, and lookup table P in $o(\log \sigma \cdot \sigma \cdot n)$ bits. A prefix sum rank query requires one superblock and block lookup as well as a constant number of arithmetic and logical operations on the last block which run all in constant time.

3 Experimental results

In this Section we will conduct computational experiments to validate our theoretical findings and to compare our FM and 2FM indices to another available implementation. All of our tests were conducted on Debian GNU/Linux 7.1 with Intel Xeon E5-2667V2 CPUs at fixed frequency of 3.3 GHz to prevent dynamic overclocking effects. All data was stored on tmpfs, a virtual file system in main memory to prevent loading data just on demand during the search and thus effecting the speed of the search by I/O operations.

In the first part of the experiments we will test FM and 2FM indices with our new data structure (EPR) in comparison to the wavelet tree (WT) implementation which was previously the generic standard implementation in SeqAn [4]. Additionally we will run the same benchmarks for other available 2FM implementations, namely the bidirectional wavelet tree by Schnattinger et al. [18] which we will call 2SCH and the balanced wavelet tree implementation with plain bit vectors and constant-time rank support in the SDSL [7] which we will refer to as 2SDSL.

The 2BWT by Lam et al. [11] is unfortunately not generic and only works for DNA alphabets. We also were not able to retrieve all hits when switching between forward and backward searches on the same pattern. Unfortunately we couldn't reach the authors and thus excluded 2BWT from our comparisons.

Another implementation that is worth mentioning is the affix array by Meyer et al. [15]. Even though the affix array implementation is generic, the construction algorithm did not terminate for alphabets other than DNA in a reasonable amount of time (several days). Unfortunately the affix array is not stand-alone but part of an application and does not provide a documented interface. Hence we were not able to include the affix array in our tests within a reasonable time frame. Meyer compares

the running time of their affix array implementation with 2SCH and states that the affix array is faster by a factor of 1.26 to 2. From that we can conclude that our 2FM index implementation using the EPR-dictionary is expected to be faster than the affix array implementation (see below).

3.1 Runtime and space consumption

For the first benchmark we want to make a comparison with alphabets of different sizes to test the predicted independence from σ for the EPR implementation. The alphabet sizes are inspired by bioinformatics applications and are of size 4 (DNA), 10 (reduced amino acid alphabet *Murphy10*), 16 (IUPAC alphabet) and 27 (protein alphabet).

We first generated a text of length 10^8 with a uniform distribution and sampled 1 million queries of length 50 from this text. The search in the FM and 2FM indices will determine the number of occurrences of the sampled strings. Our sampling will ensure that the text occurs at least once and the stepwise search is never prematurely stopped. This ensures that we have 50 million single steps in searches. The unidirectional FM indices perform backward searches while for 2FM indices we search the right half of the query first (using forward searches) and then extend the other half of the pattern to the left by backward searches.

In the following we will refer to WT and EPR as unidirectional FM indices and to 2WT and 2EPR as bidirectional FM indices, all part of the SeqAn library.

Table 1 gives an overview of the running times of all FM and 2FM index implementations. It shows the absolute runtimes as well as the speedup factor relative to the unidirectional resp. bidirectional wavelet tree implementation. WT, 2WT, 2SCH, 2SDSL are all based on wavelet trees. Our bidirectional wavelet tree implementation 2WT has a similar runtime compared to 2SDSL. It is slightly faster especially for small alphabets.

	DNA		Murphy10		IUPAC		Protein	
Index	time	factor	Time	factor	Time	factor	Time	factor
WT	6.59s	1.00	16.97s	1.00	21.61s	1.00	26.87s	1.00
EPR	**3.63s**	**1.82**	**5.35s**	**3.17**	**5.65s**	**3.83**	**6.20s**	**4.34**
2WT	9.32s	1.00	19.15s	1.00	23.44s	1.00	28.83s	1.00
2EPR	**4.69s**	**1.99**	**5.78s**	**3.31**	**5.67s**	**4.13**	**6.21s**	**4.64**
2SDSL	12.21s	0.76	20.58s	0.93	24.43s	0.96	29.76s	0.97
2SCH	14.08s	0.66	22.18s	0.86	26.11s	0.90	31.81s	0.91

Table 1: Runtimes of various implementations in seconds and their speedup factors with respect to the unidirectional wavelet tree.

Compared to the wavelet tree implementations the EPR implementation is between 80% (for DNA) and 330% (Protein) faster for unidirectional indices and between 100% (for DNA) and 360% (Protein) faster for bidirectional indices.

Since we anticipate the application of 2EPR to bioinformatics applications, we also compared the runtime of all implementations using the complete human genome

sequence. We again searched one million sampled strings of length 50 exactly as described above. The relative results were very similar to the ones in Table 1, indeed even slightly better. 2SCH crashed with this data set. 2SDSL was the slowest implementation (15.44s) followed by 2WT (1.6 times as fast) and by 2EPR (2.8 times as fast as 2SDSL) which was again the fastest implementation.

Our experiments also show that we were indeed able to eliminate the $\log \sigma$ factor of wavelet trees in practice, as predicted by Theorem 2. While the runtime for the WT implementations grows for larger alphabets with $\log \sigma$ the runtime of EPR and 2EPR increases only slightly for larger alphabets which can be explained by larger indices and therefore more cache misses. This can be seen in the following Figure in which we plot the runtime for EPR for different alphabets and the runtime of WT divided by $\log \sigma$. The resulting times develop very similarly.

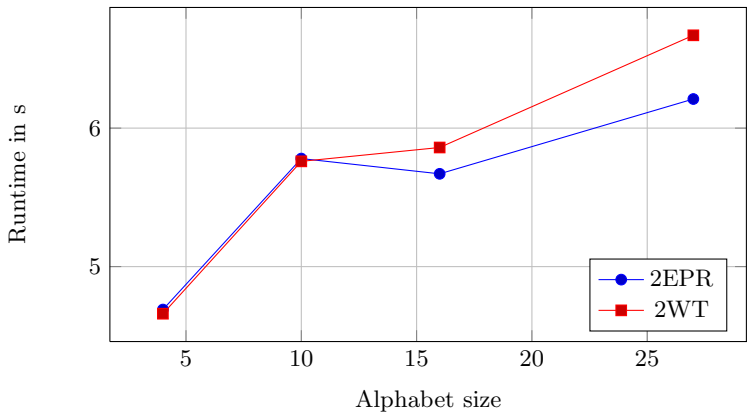

Fig. 2: Plot of the runtime for EPR for different alphabets and the runtime of WT divided by $\log \sigma$.

When we compare the runtimes of the EPR and 2EPR, they behave as expected, i.e., the unidirectional index is slightly faster, since in each step of the bidirectional index we have to synchronize two indices.

All indices implemented in SeqAn (WT, EPR, 2WT, 2EPR) support up to 3 levels for rank dictionary support: blocks, superblocks and ultrablocks. The tests presented here were performed with a 2-level rank dictionary similar to the one explained in Section 2.1 (or in the Appendix). Table 2 illustrates the practical space consumption for all previously discussed indices and of the affix array for DNA (larger alphabets did not finish within several days).

Please note, that the other implementations may use versions of rank dictionaries different to the simple one explained in Section 2.1. The numbers of FM indices given in Table 2 do neither account for storing the text itself nor for storing a compressed suffix array necessary to locate the matches in the text since the libraries use different implementations offering various space-time trade-offs. The running time of the backward and forward searches does not depend on it and the compressed suffix array

implementation is independent from the used rank dictionary and thus interchangeable. A typical compressed suffix array implementation as used in the 2SDSL takes $\frac{n}{\eta} \log n$ (when sampling on the text instead of the suffix array). For a sampling rate of 10% ($\eta = 10$) the space consumption for our experiments would be 253 MB and thus still much smaller than the affix array.

Index	DNA	Murphy10	IUPAC	Protein
EPR	42	156	227	478
2EPR	84	311	454	955
WT	30	51	60	72
2WT	60	102	120	144
2SDSL	68	105	122	145
2SCH	75	108	123	146
AF	2670	-	-	-

Table 2: Space consumption of the rank data structure in Megabyte of various implementations

The current implementation of the EPR and 2EPR in SeqAn interleaves the bit vector (i.e., the BWT) and precomputed block values but does not interleave superblock values. Reconsidering the design and storing block and superblock values close to the corresponding bit vector region could decrease the number of cache misses for one rank query to one cache miss and thus further improve the running time.

For larger alphabets one might also consider using a 3-level rank dictionary with smaller data types for blocks and superblocks which will reduce the space consumption noticeably at the expense of a slightly higher runtime (i.e., for the protein alphabet we reduce the space consumption from 955 MB to 581 MB while increasing the runtime from 6.21 to 7.67 seconds). The increased running time is due to another array lookup and thus still constant-time per step.

3.2 Effect of the low order terms for space consumption

In Table 3 we show how quickly the $o(\log \sigma \cdot \sigma \cdot n)$ data structures for rank queries can be neglected for growing n. For the WT and EPR implementations we measured the space needed for both the DNA and the IUPAC alphabet for $n = 10^4, 10^5, 10^6, 10^7, 10^8, 10^9$. We then divided the space consumption of both implementations by the factor in the $\mathcal{O}$-term, namely $\log \sigma \cdot n$.

For growing n the $\mathcal{O}$-term should dominate the low order o-term, hence we would expect the resulting number converge to a constant. This is indeed true, as can be seen in Table 3. The EPR implementation converges faster than the WT, which is expected, since our o-term is larger than the one for the WT implementations. The effect of the o-terms falls for EPR from 10^5 to 10^6 by 35 resp. 6 percent, whereas the decline for WT is steeper with 84 and 123 percent. From size 10^7 on, the low order terms are clearly dominated by the $\mathcal{O}$-terms.

Method (σ)	10^4	10^5	10^6	10^7	10^8	10^9
EPR (4)	2.4000	0.6000	0.4440	0.4292	0.4276	0.4274
EPR (16)	2.0000	1.2400	1.1680	1.1610	1.1602	1.1601
WT (4)	4.4000	0.6400	0.3480	0.3088	0.3056	0.3053
WT (16)	7.0000	0.8000	0.3580	0.3104	0.3057	0.3053

Table 3: Influence of the space consumption of the o-terms with increasing n.

4 Conclusions

In this paper we have introduced a new data structure, the EPR-dictionary, that enables constant time prefix sum computations for arbitrary, finite alphabets in $\mathcal{O}(\log \sigma \cdot n) + o(\log \sigma \cdot \sigma \cdot n)$ bits of space and works directly on the BWT. This allows two important data structures, the FM and 2FM index, to perform single search steps in time $\mathcal{O}(1)$. We implemented the dictionary in the C++ library SeqAn and used it for an implementation of an FM and 2FM index. We compared its practical performance with the previous SeqAn implementation using wavelet trees and with other openly available implementations, among them the quasi standard for succinct data structures, the SDSL. We show that the EPR-dictionary implementation supports our theoretical claims, eliminates the $\log \sigma$ factor for searching in bidirectional indices, and performs between 80% and 360% faster than the wavelet tree implementation at the expense of a higher memory consumption. We compared our 2FM implementation against the available, open implementation of Schnattinger et al. (2SCH). Our implementation is between 3 to 5.1 times faster than 2SCH and 2.6 to 4.8 faster than the 2SDSL. We also showed that the additional space consumption is easily tolerable on normal hardware.

Acknowledgments

We would like to acknowledge Enrico Siragusa for his previous implementations of the FM index in SeqAn. The first author acknowledges the support of the International Max-Planck Research School for Computational Biology and Scientific Computing (IMPRS-CBSC). We also thank Veli Mäkinen and Simon Gog for very helpful remarks on a previous version of this manuscript during the Dagstuhl seminar 16351 "Next Generation Sequencing - Algorithms, and Software For Biomedical Applications".

References

1. D. Belazzougui, F. Cunial, J. Kärkkäinen, and V. Mäkinen, *Versatile succinct representations of the bidirectional burrows-wheeler transform*, in European Symposium on Algorithms, Springer, 2013, pp. 133–144.
2. D. Belazzougui and G. Navarro, *Optimal lower and upper bounds for representing sequences*, 11 (2015), p. 31.
3. M. Burrows and D. J. Wheeler, *A block-sorting lossless data compression algorithm*, tech. report, 1994.

4. A. Döring, D. Weese, T. Rausch, and K. Reinert, *SeqAn an efficient, generic C++ library for sequence analysis*, BMC Bioinformatics, 9 (2008), p. 11, http://dx.doi.org/10.1186/1471-2105-9-11.
5. P. Ferragina and G. Manzini, *Opportunistic data structures with applications*, ANNUAL SYMPOSIUM ON FOUNDATIONS OF COMPUTER SCIENCE, (2000), http://dx.doi.org/10.1109/SFCS.2000.892127.
6. P. Ferragina, G. Manzini, V. Mäkinen, and G. Navarro, *Compressed representations of sequences and full-text indexes*, ACM Transactions on Algorithms (TALG), 3 (2007), p. 20.
7. S. Gog, T. Beller, A. Moffat, and M. Petri, *From theory to practice: Plug and play with succinct data structures*, in 13th International Symposium on Experimental Algorithms, (SEA 2014), 2014, pp. 326–337.
8. R. Grossi, A. Gupta, and J. Vitter, *High-order entropy-compressed text indexes*, Proceedings of the fourteenth annual ACM- , (2003).
9. G. J. Jacobson, *Succinct static data structures*, (1988).
10. T. Lam, R. Li, A. Tam, S. Wong, and E. Wu, *High throughput short read alignment via bi-directional BWT*, in Proceedings of BIBM, 2009, pp. 31–36, http://dx.doi.org/10.1109/BIBM.2009.42.
11. T. Lam, W. Sung, S. Tam, C. Wong, and S. Yiu, *Compressed indexing and local alignment of DNA*, 24 (2008), pp. 791–797, http://dx.doi.org/10.1093/bioinformatics/btn032.
12. B. Langmead and S. L. Salzberg, *Fast gapped-read alignment with Bowtie 2*, Nat. Methods, 9 (2012), pp. 357–359.
13. H. Li, *Aligning sequence reads, clone sequences and assembly contigs with BWA-MEM*, (2013).
14. H. Li and R. Durbin, *Fast and accurate short read alignment with burrows-wheeler transform*, Bioinformatics, 25 (2009), pp. 1754–1760, http://dx.doi.org/10.1093/bioinformatics/btp324.
15. F. Meyer, S. Kurtz, R. Backofen, S. Will, and M. Beckstette, *Structator: fast index-based search for RNA sequence-structure patterns.*, BMC Bioinformatics, 12 (2011), p. 214, http://dx.doi.org/10.1186/1471-2105-12-214.
16. G. Navarro and E. Providel, *Fast, small, simple rank/select on bitmaps*, International Symposium on Experimental Algorithms, (2012), http://dx.doi.org/10.1007/978-3-642-30850-5_26.
17. M. Santiago, M. Sammeth, R. Guigo, and P. Ribeca, *The GEM mapper: fast, accurate and versatile alignment by filtration.*, Nature Methods, 9 (2012), pp. 1185–1188, http://dx.doi.org/10.1038/nmeth.2221.
18. T. Schnattinger, E. Ohlebusch, and S. Gog, *Bidirectional search in a string with wavelet trees and bidirectional matching statistics*, Information and Computation, 213 (2012), pp. 13–22, http://dx.doi.org/10.1016/j.ic.2011.03.007.
19. E. Siragusa, *Approximate string matching for high-throughput sequencing*, PhD thesis, Freie Universitt Berlin, 2015.
20. E. Siragusa, D. Weese, and K. Reinert, *Fast and accurate read mapping with approximate seeds and multiple backtracking*, Nucleic Acids Res, 41 (2013), pp. e78–e78, http://dx.doi.org/10.1093/nar/gkt005.

Appendix

In the appendix we give for the reader not familiar with FM and 2FM indices a short introduction.

Introduction to the FM and 2FM Index

Given a text T of length n over an ordered, finite alphabet $\Sigma = \{c_1, \ldots, c_\sigma\}$ with $\forall\, 1 \leq i < \sigma : c_i <_{lex} c_{i+1}$, let $T[i]$ denote the character at position i in T, $\cdot$ the concatenation operator and $T[1..i]$ the prefix of T up to the character at position i. T^{rev} represents the reversed text. We assume that T ends with a sentinel character $\$ \notin \Sigma$ that does not occur in any other position in T and is lexicographically smaller than any character in Σ. The FM index needs the Burrows-Wheeler transform (BWT) of T. The BWT is the concatenation of characters in the last column of all lexicographically sorted cyclic permutations of the string T (see Figure 3 for an example). We will refer to the BWT as L.

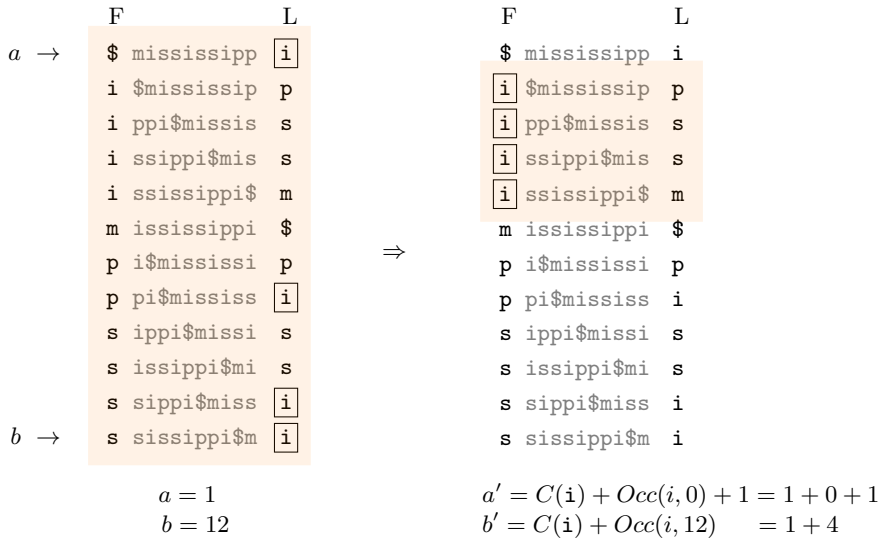

Fig. 3: First step of the backwards search for $P =$ `ssi` in the FM-index for the text $T =$ `mississippi$`. The first interval $[a, b]$ is the whole range $[1, 12]$. From all matrix rows we search those beginning with the last pattern character $P[3] = i$. From $Occ(i, 1) = 0$ and $Occ(i, 12) = 4$ follows $a' = C(i) + 0 + 1 = 2$ and $b' = C(i) + 4 = 5$.

In contrast to suffix trees or suffix arrays, where a prefix P of a pattern is extended by characters to the right (referred to as forward search $P \to Pc$ for $c \in \Sigma$), the FM index can only be searched using backward search, i.e., extending a suffix P' by characters to the left, $P' \to cP'$. Performing a single character backward search of c in the FM index will require two pieces of information. First, $C(c)$, the number of characters in L that are lexicographically smaller than c, second, $Occ(c, i)$, the

number of c's in $L[1..i]$. Given a range $[a, b]$ for P; i.e., the range in the sorted list of cyclic permutations that starts with P, we can compute the range $[a', b']$ for cP as follows: $[a', b'] = [C(c) + Occ(c, a-1) + 1, C(c) + Occ(c, b)]$. We will refer to the BWT together with tables C and Occ as *FM index* $\mathcal{I}$ (see Figure 3 for an example of one search step).

The 2FM index maintains two FM indices $\mathcal{I}$ and $\mathcal{I}^{rev}$, one for the original text T and one for the reversed text T^{rev}. Searching a pattern left to right on the original text (i.e., conducting a forward search) corresponds to a backward search in $\mathcal{I}^{rev}$; searching a pattern right to left in the original text corresponds to a backward search in $\mathcal{I}$. The difficulty is to keep both indices *synchronized* whenever a search step is performed. W.l.o.g. we assume that we want to extend the pattern to the right, i.e., perform a forward search $P \to Pc_j$ for some character c_j. First, the backward search $P^{rev} \to c_j P^{rev}$ is carried out on $\mathcal{I}^{rev}$ and its new range $[a'_{rev}, b'_{rev}] = [C(c_j) + Occ(c_j, a_{rev} - 1) + 1, C(c_j) + Occ(c_j, b_{rev})]$ is computed. The new range in $\mathcal{I}$ can be calculated using the interval $[a, b]$ for P in $\mathcal{I}$ and the range size of the reversed texts index $[a', b'] = [a + smaller, a + smaller + b'_{rev} - a'_{rev}]$. To compute $smaller$, Lam et al. [10] propose to perform $\mathcal{O}(\sigma)$ backward searches $P^{rev} \to c_i P^{rev}$ for all $1 \leq i < j$ and sum up the range sizes, i.e., $smaller = \sum_{1 \leq i < j} Occ(c_i, b_{rev}) - \sum_{1 \leq i < j} Occ(c_i, a_{rev} - 1)$ leading to a total running time of $\mathcal{O}(\sigma)$ (See Figure 4 for an illustration).

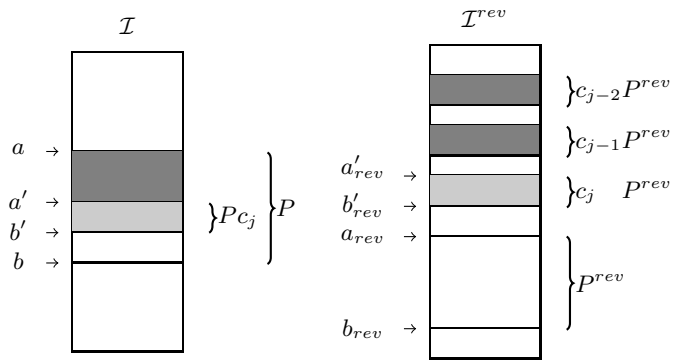

Fig. 4: When conducting a forward search $P \Rightarrow Pc_j$ we need to determine the subinterval of the suffix array interval for P which is depicted on the left. In order to determine the start, we can compute in $\mathcal{I}^{rev}$ the size of the intervals for all characters smaller then c_j, depicted in dark gray on the right. The sum of all those sizes is exactly the needed offset from the beginning of the interval for P in $\mathcal{I}$.

The implementation of the occurrence table Occ is usually *not* done by storing explicitly the values of the entire table. Instead of storing the entire $Occ : \Sigma \times \{1, \ldots, n\} \to \{1, \ldots, n\}$ one uses the more space-efficient *constant time rank dictionary*: for every $c \in \Sigma$ a bit vector $B_c[1..n]$ is constructed such that $B_c[i] = 1$ if and only if $L[i] = c$. Thus the occurrence value equals the number of 1's in $B_c[1..i]$, i.e., $Occ(c, i) = rank(B_c, i)$. Jacobson [9] showed that rank queries can be answered in constant time using only $o(n)$ additional space per bit vector. Since then many other constant time rank query data structures have been proposed. For an overview

we refer the reader to [16] containing a comparison of various implementations. Since we will make also use of this technique, we explain the most simple idea, namely the one for *2-level rank dictionaries* in the following paragraph.

Constant time rank queries

In order to store partial prefix sums, the technique uses two levels of lookup table, called *blocks* and *superblocks*. Given a bit vector B of length n we divide it into blocks of length ℓ and superblocks of length ℓ^2 where

$$\ell = \left\lfloor \frac{\log n}{2} \right\rfloor .$$

For both, blocks and superblocks we allocate arrays M and M' of sizes $\left\lfloor \frac{n}{\ell} \right\rfloor$ and $\left\lfloor \frac{n}{\ell^2} \right\rfloor$ respectively (see Figure 5 for an illustration).

For the m-th superblock we store the number of 1's from the beginning of B to the end of the superblock in $M'[m] = rank(B, m \cdot \ell^2)$. As there are $\left\lfloor \frac{n}{\ell^2} \right\rfloor$ superblocks, M' can be stored in $\mathcal{O}\left(\frac{n}{\ell^2} \cdot \log n\right) = \mathcal{O}\left(\frac{n}{\log n}\right) = o(n)$ bits. For the m-th block we store the number of 1's from the beginning of the overlapping superblock to the end of the block in $M[m] = rank\big(B[1 + k\ell..n], (m-k) \cdot \ell\big)$, where $k = \left\lfloor \frac{m-1}{\ell} \right\rfloor \ell$ is the total number of blocks in all superblocks left of the current superblock. M has $\left\lfloor \frac{n}{\ell} \right\rfloor$ entries and can be stored in $\mathcal{O}\left(\frac{n}{\ell} \cdot \log \ell^2\right) = \mathcal{O}\left(\frac{n \cdot \log\log n}{\log n}\right) = o(n)$ bits.

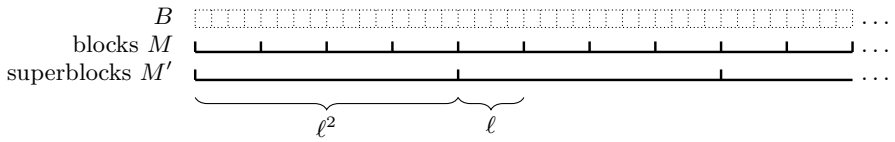

Fig. 5: 2-level dictionary. Blocks and superblocks are allocated for each character (only one shown).

Given a rank query $rank(B, i)$, one can now add the corresponding superblock and block values. But we still have to account for the 1's in the block covering position i (in case i is not at the end of a block). Let P be a precomputed lookup table such that for each possible bit vector V of length ℓ and $i \in [1..\ell]$ holds $P[V][i] = rank(V, i)$. V has $2^\ell \cdot \ell$ entries of values at most ℓ and thus can be stored in

$$\mathcal{O}\left(2^\ell \cdot \ell \cdot \log \ell\right) = \mathcal{O}\left(2^{\frac{\log n}{2}} \cdot \log n \cdot \log\log n\right) = \mathcal{O}\left(\sqrt{n} \cdot \log n \cdot \log\log n\right) = o(n)$$

bits. We now decompose a rank query into 3 subqueries using the precomputed tables. For a position i we determine the index $p = \left\lfloor \frac{i-1}{\ell} \right\rfloor$ of next block left of i and the index $q = \left\lfloor \frac{p-1}{\ell} \right\rfloor$ of the next superblock left of block p. Then it holds:

$$rank(B, i) = M'[q] + M[p] + P\big[B[1 + p\ell..(p+1)\ell]\big]\big[i - p\ell\big].$$

Since the text T of length n has to be addressed, we assume that a register has at least size $\lceil \log n \rceil$. Thus $B[1 + p\ell..(p + 1)\ell]$ fits into a single register and can be determined in $\mathcal{O}(1)$ time. Therefore a rank query can be answered in $\mathcal{O}(1)$ time. In practice the precomputed lookup table P is replaced by a popcount operation on the CPU register and we have only two lookup operations.

One can now replace the occurrence table by this 2-level dictionary, i.e., by creating a bit vector for every $c \in \Sigma$ and setting it to 1 if the character occurs in the BWT L. This results in $\mathcal{O}(\sigma \cdot n) + o(\sigma \cdot n)$ bits space requirement.